arXiv:1512.02271v1 [math.OC] 7 Dec 2015

Optimal strategies for the control of autonomous vehicles in data assimilation

D. McDougall[a,*], R.O. Moore[b]

[a] *201 E. 24th St., Stop C0200*
Institute for Computational Engineering and Sciences
University of Texas at Austin
Austin, TX, 78712
[b] *Department of Mathematical Sciences*
New Jersey Institute of Technology
University Heights, Newark, NJ 07102

Abstract

We propose a method to compute optimal control paths for autonomous vehicles deployed for the purpose of inferring a velocity field. In addition to being advected by the flow, the vehicles are able to effect a fixed relative speed with arbitrary control over direction. It is this direction that is used as the basis for the locally optimal control algorithm presented here, with objective formed from the variance trace of the expected posterior distribution. We present results for linear flows near hyperbolic fixed points.

Keywords: Bayesian inverse problem, Lagrangian data assimilation, Optimal control, Ocean glider
2010 MSC: 49M, 62F, 62L, 93C, 65C

1. Introduction

The need for a more accurate and better resolved estimate of oceanic flows is being driven by a number of pressing global issues, including the crisis facing many species of fish and waterborne organisms, the mitigation of pollutants resulting from spills and offshore contamination, and the important role played by ocean dynamics on climate change. Scientific efforts to estimate ocean flow began in the 1980s with the work of Robinson [1], but has enjoyed limited success due to a lack of observational data. In an effort to improve the current state of understanding of the world's oceans, autonomous vehicles (AVs) are being deployed for the collection of physical oceanography data in a growing number of projects around the globe. One example of AVs are autonomous

*Corresponding author
Email addresses: damon@ices.utexas.edu (D. McDougall), rmoore@njit.edu (R.O. Moore)

underwater vehicles (AUVs), which are equipped with adjustable wings that convert vertical momentum induced by battery-powered changes in buoyancy into forward momentum.

Notwithstanding a myriad of challenges caused by their immersion in a high-inertia environment and the limited power and bandwidth available for communication, the effectiveness of AVs comes from their ability to remain in the field for prolonged deployment and from their capacity for controlled self-propulsion. To understand why this control is so important, one need simply consult the growing literature on filtering of complex systems. The dimension associated with a meaningful estimate of an oceanic velocity field is far larger than a typical AV cohort size, so that the observations are necessarily sparse in space. They may also be sparse in time for autonomous underwater vehicles (AUVs) that are unable to communicate while submerged. Even large cohorts of Lagrangian drifters that simply advect with the flow are often subject to collapse onto a small subset of stable dynamic features of the flow, rendering their observations increasingly less informative with time. Recent demonstrations that such drifters face inherent information barriers suggest that simply increasing the size of a deployed host is not effective without also implementing a method to allow them to cross flow barriers [2].

To better understand this issue, we reconstruct a two-dimensional flow field from observations taken by AVs that are capable of limited locomotion in addition to advecting with the flow. We refer to these AVs as *gliders*, in distinction to drifters. They are capable of measuring the surrounding fluid velocity either directly using Doppler velocimetry or indirectly using sequential position measurements, depending on the instrumentation on board. In this work we focus on the former case and use the velocity observations to perform an "identical twin" experiment based on a perfect model, with the goal of constructing an estimate of the true flow field function from observations, along with an associated uncertainty. Our approach to control follows the underlying philosophy of [3] that mesoscale ocean flow fields are dominated by coherent features, such as jets and eddies, which not only dictate the sites where the most informative observations can be taken to estimate these structures, but also provide the most effective transport mechanisms for navigation of the physical domain.

If the ocean drifter can be controlled to move into and through these structures then the information gained should be richer in terms of capturing the key properties of the flow field. Whereas in [3] this was shown by utilizing an ad hoc control strategy, the present work attempts to put this philosophy on a more systematic footing through the application of an *optimal* control strategy. Each optimal control calculation is inserted between successive observations in a sequential filter, utilizing the posterior field estimate to computing a control that most effectively minimizes the next expected posterior variance.

In general, reconstructing the flow requires one to solve an inverse problem which is most naturally posed in a Bayesian framework ([4]), the solution to which is a probability distribution on the appropriate function space or parameter space, depending on the velocity field's representation. Kalman filters and Kalman smoothers ([5, 6, 7, 8, 9, 10]) completely quantify this distribution

by its first two moments in the case of linear processes and Gaussian initial distributions, and are popular approximate methods for nonlinear flows and non-Gaussian distributions. The Kalman filter is the assimilation mechanism used here, ostensibly due to the use of linear stationary velocity fields and Gaussian distributions but also to provide the straightforward extension to weakly nonlinear flow models through the extended Kalman filter [4].

Alternative approaches include variational methods that view the solution not as a distribution but as the argument of a cost function that is optimized [11, 12, 13, 14, 15, 16, 17, 18, 19], and particle filtering methods ([20, 21]) that approximate the continuous posterior random variable by a discrete set of state realizations (particles) with associated weights. Updating the particles and weights as new observations are made is difficult and can lead to degenerate posteriors in high-dimensional problems [22]. Sampling methods utilize the Metropolis-Hastings algorithm ([23, 24]) to sequentially generate correlated samples from the posterior distribution ([25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40]). Sampling methods are ideal for the online computation of moments through the use of unbiased estimators, but are not suited to applications where the data arrive sequentially, and are expensive in settings where the distribution must be used in an inverse problem to determine the minimizing control.

Section 2 describes the setup of the problem under consideration in a relatively broad setting, including the basic flow assumptions and assimilation model, followed by a derivation of the method used to find local minimizers of the objective. Section 3 identifies an existence problem with the variational formulation and introduces a simple approach to address it. Section 4 provides numerical results for four linear time-independent velocity fields modeling flow near hyperbolic fixed points with different stability properties. Section 5 offers conclusions and discusses extensions to the methodology.

2. Setup

The context of this study is oceanographic data assimilation, where the goal is to estimate a velocity field conditioned on sparse observational data recorded by Lagrangian *gliders* whose positions evolve according to the flow being assimilated in addition to a modest capacity for self-propulsion. The gliders therefore move according to

$$\dot{z} = v(z(t), t) + u(t), \tag{1}$$

with $z \in \mathbb{R}^{dK}$ contains the positions of K gliders in a d-dimensional domain (here, we take $d = 2$), and the flow evolves according to a model expressed as

$$\frac{\partial v}{\partial t} = f(v). \tag{2}$$

The glider's self-propulsion is expressed by control $u(t)$, where $u(t) \equiv 0$ represents the uncontrolled case (i.e., Lagrangian drifters [41]) and $|u(t)| \equiv u_{\max}$

represents the case considered here with the relative speed of gliders fixed at $u_{\max}$ [42].

The direct velocity observations considered here take the form

$$y_i = v(z(t_i), t_i) + \eta_i, \tag{3}$$

where $\eta_i \sim \mathcal{N}(0, R_i)$ is i.i.d Gaussian-distributed measurement noise. Here, $y_i \in \mathbb{R}^{dK}$ and $v(z,t) := (v(z_1,t), v(z_2,t), \ldots, v(z_K,t))^T$. The mapping between "state space" (estimates of $v(x,t)$) and "observation space" (values of y) is referred to as the observation operator which, in this case, is formally given by the convolution

$$H(v) = \delta(x - z(t_i)) * v. \tag{4}$$

If the observational data consist of glider positions, then an alternative to inferring velocities through, for example, dead reckoning before assimilation is to augment the state vector to include the glider positions [43]. If the mean discretized velocity state estimate $\hat{v}$ is dN-dimensional, this produces an augmented state vector $\tilde{v} := (\hat{z}^T, \hat{v}^T)^T$ of dimension $d(N+K)$, with the advantage that the observation operator is simply

$$H(\tilde{v}) = \begin{pmatrix} I_{dK \times dK} & O_{dK \times NK} \end{pmatrix} \begin{pmatrix} \hat{z} \\ \hat{v} \end{pmatrix}. \tag{5}$$

2.1. Assimilation

There exist a wide variety of techniques for the assimilation of Lagrangian observational data into a velocity model, depending on the context. For linearly evolving flows, filters based on Gaussian or conditionally Gaussian [44] processes such as the Kalman filter are appropriate, with suitable modification to reconcile Eulerian and Lagrangian descriptions [45, 43]. These methods have been extended with some success to weakly nonlinear flows; fully nonlinear flows require alternative methods such as particle filters or variational frameworks to explore the posterior distribution.

The ultimate goal is to combine both the flow model in Eqn. 2 and the observations in Eqn. 3 to obtain an estimate (with uncertainty) of the underlying flow v. In terms of the Bayesian formulation, we seek the posterior probability distribution of the flow field conditioned on noisy observations, i.e.,

$$\frac{\mu^y}{\mu_0}(v) \propto \mathbb{P}(y|v), \tag{6}$$

where the posterior distribution $\mathbb{P}(v|y)$ has measure μ^y and the prior distribution $\mathbb{P}(v)$ has measure μ_0. This formulation of Bayes's rule, in terms of the Radon-Nikodym derivative, is necessary to ensure well-posedness of the posterior distribution on function space [46]. In the case of a linear evolution model f, an initially Gaussian velocity field estimate maintains this property, allowing its complete characterization through two moments, the mean $\hat{v}$ and covariance

$\hat{P}$, with

$$\hat{v}_i = \mathbb{E}(v|y_i) \tag{7}$$

$$\hat{P}_i(x,x') = \mathbb{E}((v(x) - \mathbb{E}(v(x)))(v(x') - \mathbb{E}(v(x')))^\top). \tag{8}$$

Here, $\hat{v}_i$ and $\hat{P}_i$ incorporate observations up to time t_i. They are updated sequentially upon the arrival of new data. If observations take the form expressed in Eqn. 4 the update formula is given by the Kalman filter equations ([47]),

$$\hat{v}_i = \hat{v}_{i-1} + K_i(y_i - \hat{v}_{i-1}(z(t_i))), \tag{9}$$

$$\hat{P}_i = \hat{P}_{i-1} + K_i H(v;u)\hat{P}_{i-1}. \tag{10}$$

Here K_i is called the Kalman gain and is given by,

$$K_i = \hat{P}_{i-1}H_i^\top(H_i\hat{P}_{i-1}H_i^\top + R)^{-1}. \tag{11}$$

2.2. Controlling the drifters

Our objective is to control gliders to where their measurements of the flow will minimize the trace of the posterior covariance,

$$\Gamma(z(t_i)) = \operatorname{tr}\left(\hat{P}_{i-1} - \hat{P}_{i-1}H_i^\top(H_i\hat{P}_{i-1}H_i^\top + R)^{-1}H_i\hat{P}_{i-1}\right). \tag{12}$$

The trace of the posterior covariance matrix can be thought of as the 'global' uncertainty. Therefore minimising this quantity can be thought of as maximising our knowledge of the underlying ocean flow, which is crucial for use in a prediction.

To impose the fixed relative speed constraint, we take a similar approach to Zermelo's problem [48] and replace the $2K$-dimensional control u with

$$u = u_{\max}\hat{u} := u_{\max}(\cos\theta_1, \sin\theta_1, \cos\theta_2, \ldots, \sin\theta_K)^\top, \tag{13}$$

reducing the number of degrees of freedom to K angles θ_k. The objective function we aim to minimize is therefore

$$\Phi(z(t_j)) = -\Gamma(z(t_j)) - \int_{t_{j-1}}^{t_j} \lambda(t)^\top(\dot{z} - v(z) - u_{\max}\hat{u}(t))\,dt. \tag{14}$$

Setting the first variation of Φ with respect to $z(t)$ equal to zero gives the Euler-Lagrange equations,

$$\dot{z} = v(z) + u_{\max}\hat{u}(t), \quad z(t_{j-1}) = z_{j-1}, \tag{15}$$

$$\dot{\lambda} = -\left(\frac{\partial v}{\partial z}\right)^\top \lambda, \quad \lambda(t_j) = -\left(\left.\frac{\partial \Gamma}{\partial z}\right|_{t_j}\right)^\top. \tag{16}$$

The costate λ and control $\hat{u}$ are related through

$$\lambda(t_j)^\top \frac{\partial \hat{u}}{\partial \theta} = 0, \tag{17}$$

where $\partial\hat{u}/\partial\theta$ is a block-diagonal matrix of 2-vectors $(-\sin\theta_k, \cos\theta_k)^\top$, with size $2K \times K$. Since each of these 2-vectors is orthogonal to the corresponding 2-vector component of $\hat{u}$, this condition is equivalent to $2K$-vector λ taking the form

$$\lambda = -\begin{pmatrix} \tilde{\lambda}_1 \hat{u}_1 \\ \tilde{\lambda}_2 \hat{u}_2 \\ \vdots \end{pmatrix}, \tag{18}$$

essentially representing λ as a vector of polar forms.

The above Euler-Lagrange system of state and costate equations naturally suggests a shooting method to compute the minimizer of Φ. Here we take an alternative approach and derive a relaxation method that is based on the expression of the coupled one-way initial- and terminal-value problems for the state and costate. We arrive at a single second-order boundary value problem that we reformulate as the steady state of a parabolic PDE. Specifically, differentiating Eqn. 16 and using Eqn. 18, we arrive at the equivalent equation for K-vector θ:

$$\dot{\theta} = -\operatorname{diag}\hat{u}^\top \operatorname{diag} J^\top \left(\frac{\partial v}{\partial z}\right)^\top \hat{u}, \tag{19}$$

where $\operatorname{diag}\hat{u}$ is the $2K \times K$ matrix with each 2-vector u_k along the diagonal, and

$$J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \tag{20}$$

Note that the K costate amplitudes $\tilde{\lambda}_k$ have dropped out, leaving just the angles θ_k. These angles are decoupled, i.e., Eqn. 19 is just a K-fold copy of

$$\dot{\theta}_k = -\begin{pmatrix} -\sin\theta_k & \cos\theta_k \end{pmatrix} \begin{pmatrix} \frac{\partial v_{kx}}{\partial x_k} & \frac{\partial v_{ky}}{\partial x_k} \\ \frac{\partial v_{kx}}{\partial y_k} & \frac{\partial v_{ky}}{\partial y_k} \end{pmatrix} \begin{pmatrix} \cos\theta_k \\ \sin\theta_k \end{pmatrix}. \tag{21}$$

The (forward) state equations and (backward) costate equations are also decoupled, in addition to the boundary conditions for the state at t_{j-1}, leaving as the only source of coupling the terminal condition for λ and therefore θ. This strongly suggests that an iterative numerical method using decoupled path optimization should be possible. We will proceed with the formulation of the relaxation method for a single AUV, with the understanding that all AUV paths are coupled at t_j.

Differentiating Eqn. 16 and making use of Eqn. 19, we arrive at the boundary value problem

$$\ddot{z} = \frac{\partial v}{\partial z}\dot{z} - \frac{1}{u_{\max}^2} J(\dot{z} - v)(\dot{z} - v)^\top J^\top \left(\frac{\partial v}{\partial z}\right)^\top (\dot{z} - v). \tag{22}$$

The boundary condition at initial time t_j is linear and Dirichlet:

$$z(t_{j-1}) = z_{j-1}. \tag{23}$$

At the terminal time t_j, however, the boundary condition is nonlinear unless v is linear and Γ is quadratic:

$$\dot{z}(t_j) = v(z(t_j)) + u_{\max} \left. \frac{(\partial\Gamma/\partial z)^\top}{\|\partial\Gamma/\partial z\|} \right|_{z(t_j)}. \tag{24}$$

As noted earlier, this boundary condition also couples all glider positions at t_j through the posterior covariance trace $\Gamma(z_1, \ldots, z_K)$.

2.2.1. Numerical method

Equation 22 with boundary conditions 23 and 24 expresses a two-point boundary value problem for which a number of solution techniques exist [49]. We adopt a relaxation approach to finding solutions, in which we introduce an artificial time τ and formulate a partial differential equation (PDE) with solutions to the two-point boundary problem as its steady states:

$$\frac{\partial z}{\partial \tau} = \ddot{z} - \frac{\partial v}{\partial z}\dot{z} + \frac{1}{u_{\max}^2} J(\dot{z} - v)(\dot{z} - v)^\top J^\top \left(\frac{\partial v}{\partial z}\right)^\top (\dot{z} - v). \tag{25}$$

This PDE is solved with boundary conditions as above and with a simple initial condition such as the uniform state $z(t) \equiv z_{j-1}$. The PDE is integrated in τ until the residual of Eqn. 22, i.e., $\|\partial z/\partial\tau\|$, drops below a threshold value. We use central differences in t and a semi-implicit method in τ, with $\ddot{z}$ computed implicitly and the remaining terms on the right-hand side of Eqn. 25 computed explicitly.

The minimization problem posed above is only expected to have unique solutions under very specific conditions, such as linear velocity and quadratic objective, so the initial condition chosen for the relaxation method will play a large role in determining the local minimum found. This is discussed in Sec. 5.

As noted earlier, these equations are uncoupled apart from the boundary condition at t_j. We therefore adopt the following iterative method defined on a uniform partition of $(t_{j-1}, t_j]$:

1. Initialize all paths $z_k(t, 0)$, e.g., using constants $z_k(t, 0) \equiv z_{k,j-1}$ or straight-line paths connecting $z_{k,j-1}$ to nearby maxima of Γ.
2. For each AUV $k = 1, \ldots, K$, update the path by integrating Eqn. 25 until the residual (i.e., right-hand side) has dropped below a prescribed threshold, using Γ evaluated at fixed z_l, $l = 1, \ldots, k-1, k+1, \ldots, K$ (using the most recently computed path iterate available).
3. Repeat 2 until all updated residuals are below threshold.

3. Convexity and regularization

It is clear that convexity for this optimization problem depends entirely on the nature of the velocity field $v(z)$ and cost function $\Gamma(z(t_j))$. Nontrivial v and Γ are expected to have multiple local minimizers and, more critically, are not

guaranteed to admit a solution within the class of admissible trajectories $z(t)$. The focus of this work is simply to find a reasonable local minimizer as dictated by the choice of initial condition for the relaxation method, so the discussion of minimizer selection is deferred to future work. It is possible to find scenarios even in the case of constant velocity and quadratic cost where a solution fails to exist, which our numerical studies corroborate; we analyse this scenario below to motivate a regularization of the terminal condition 24.

If $v \equiv v_0$ is constant, relaxation method 25 collapses to the simple heat equation,

$$z_\tau = \ddot{z}, \quad z(t_0) = z_0, \tag{26}$$

where assuming $\Gamma(z) = -\frac{1}{2} z^T B z$ gives

$$\dot{z}(t_1) = v_0 - u_{\max} \frac{(B + B^T) z(t_1)}{\|(B + B^T) z(t_1)\|}. \tag{27}$$

Coupling between the gliders is quantified by off-diagonal elements, i.e., elements outside the K 2×2 submatrices that lie along the main diagonal. To motivate the regularization below it is sufficient to consider decoupled gliders with $B = I$, although the same analysis is possible for any B that is symmetric and positive-definite. The boundary condition 27 is then simply

$$\dot{z}(t_1) = v_0 - u_{\max} \frac{z(t_1)}{\|z(t_1)\|} \tag{28}$$

and steady states of Eqn. 26 must satisfy

$$z(t_1) = z(t_0) + (t_1 - t_0) \left(v_0 - u_{\max} \frac{z(t_1)}{\|z(t_1)\|} \right). \tag{29}$$

Figure 1 illustrates that the existence and uniqueness of solutions in this case depends entirely on the difference between the uncontrolled distance from the origin (i.e., the global maximum of Γ) at time t_1 and the maximum relative distance possible under control $u_{\max}$, i.e.,

$$\delta L = \|z(t_0) + (t_1 - t_0) v_0\| - (t_1 - t_0) u_{\max}. \tag{30}$$

If $\delta L > 0$, as in Fig. 1a, a solution exists and is unique. If $\delta L < 0$, Fig. 1b demonstrates that no solution exists. If $\delta L = 0$, the obvious solution has the glider land exactly at the maximum of Γ, although the boundary condition 27 is not well-defined here.

The failure to find a solution is directly related to the fixed-speed constraint on each glider; if its speed is large enough to send it past a nearby local maximum of the objective, this implies the lack of a minimizer in the space of smooth functions. This problem can be addressed either by including piecewise smooth curves in the class of admissible paths or by replacing the fixed-speed constraint with a speed penalty in the objective function. To minimize the impact on paths not affected by the problem discussed here, we adopt an intermediate

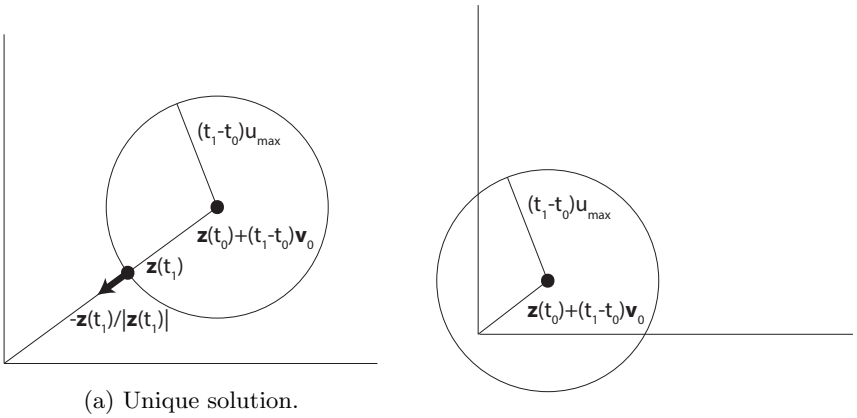

(a) Unique solution.

(b) No solution.

Figure 1: (A) Illustration of unique solution to condition 29 when $\|z(t_0) + (t_1 - t_0)v_0\| > (t_1 - t_0)u_{\max}$. (B) Illustration that no solution exists when $\|z(t_0) + (t_1 - t_0)v_0\| < (t_1 - t_0)u_{\max}$.

approach of altering boundary condition 24 to allow violations of the fixed-speed constraint for gliders close to a local maximum, where the gradient of the objective drops below a specified magnitude. The regularized boundary condition is as follows:

$$\dot{z}(t_j) = v(z(t_j)) + \begin{cases} u_{\max} \left. \frac{(\partial\Gamma/\partial z)^\top}{\|\partial\Gamma/\partial z\|} \right|_{z(t_j)} & \text{if } \|\partial\Gamma/\partial z\| > \gamma, \\ \frac{u_{\max}}{\gamma} \left. (\partial\Gamma/\partial z)^\top \right|_{z(t_j)} & \text{otherwise.} \end{cases} \tag{31}$$

The impact of this regularization on the present case is that condition 29 is replaced by

$$z(t_1) = z(t_0) + (t_1 - t_0)\left(v_0 - \frac{u_{\max}}{\gamma} z(t_1)\right), \tag{32}$$

which has solution

$$z(t_1) = \frac{\gamma}{\gamma + (t_1 - t_0)u_{\max}} \left(z(t_0) + (t_1 - t_0)v_0\right). \tag{33}$$

This solution is only relevant if $\|\partial\Gamma/\partial z\| < \gamma$ at this value of z, i.e., if $\gamma > \delta L$, which is always true in the case where no solution exists in the absence of regularization. This solution limits to $z(t_1) = 0$ (i.e., to the global maximum of Γ) as $\gamma \to 0$ and to $z(t_0) + (t_1 - t_0)v_0$ (i.e., to the uncontrolled case) as $\gamma \to \infty$.

In practice, the value of γ is chosen so to be small enough that it has no impact on the majority of trajectories that fall into the category of Fig. 1a, but large enough that it successfully resolves all cases suffering from a lack of existence due to the scenario depicted in Fig. 1b.

4. Results

To test our method we consider a stationary linear flow

$$v(z) = v_0 + Ax, \quad \text{with} \quad A \in \mathbb{R}^{2\times 2}.$$

The state vector $(v_{0x}, v_{0y}, A_{11}, A_{12}, A_{21}, A_{22})^T$ is 6-dimensional with observation operator

$$H_e = \begin{pmatrix} 1 & 0 & x_1 & y_1 & 0 & 0 \\ 0 & 1 & 0 & 0 & x_1 & y_1 \\ & & \vdots & & & \\ 1 & 0 & x_K & y_K & 0 & 0 \\ 0 & 1 & 0 & 0 & x_K & y_K \end{pmatrix}, \tag{34}$$

mapping from $\mathbb{R}^6$ to the set of $2K$ glider velocities. Numerical simulations were performed to infer this linear velocity field using four different "ground truths". In each case, the (unknown) constant component was given by $v_0 = (1/2,\, -1/2)^T$, but the velocity's Jacobian A was chosen to produce (a) a center, (b) an unstable node, (c) a saddle, and (d) a stable node. Each case is considered below, where the numerical simulations were generated with cohorts of 1, 2, 5, and 10 gliders, each taking 100 sequential velocity observations with time interval of 0.1 units between successive observations. The estimate was initialized at zero with variance of 10^6. Observational noise was set at unit variance. The maximum glider speed was set to $u_{\max} = 1$. In order to assess performance, all runs were performed with optimal control obtained using Eqn. 25, with no control, and with a fixed control for each assimilation stage with randomly selected direction. Computations of the optimal control that failed to converge were redone with the regularization discussed in Sec. 3 using

$$\gamma = u_{\max}(t_j - t_{j-1}) \| \frac{\partial^2 \Gamma}{\partial z^2} \|. \tag{35}$$

4.1. Flow about center

The simulations performed around a center use a velocity field given by

$$v_0 = \begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, \qquad A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \tag{36}$$

such that the center is located at $(-1/2, -1/2)$. Figure 2 demonstrates the advantage of using optimal control, represented by solid lines, over no control, represented by dashed lines, for cohort sizes of 1, 2, 5, and 10 in terms of the covariance trace and the root-mean-square (RMS) error. All cases show systematic improvement with the use of optimal control, with this improvement varying according to the usable information contained in the velocity field estimate. The Kalman filter estimates have necessarily decreasing covariance trace while the RMS error shows an initial increase as the first few assimilations place too much confidence on random point estimates. In the 1- and 2-glider cohorts,

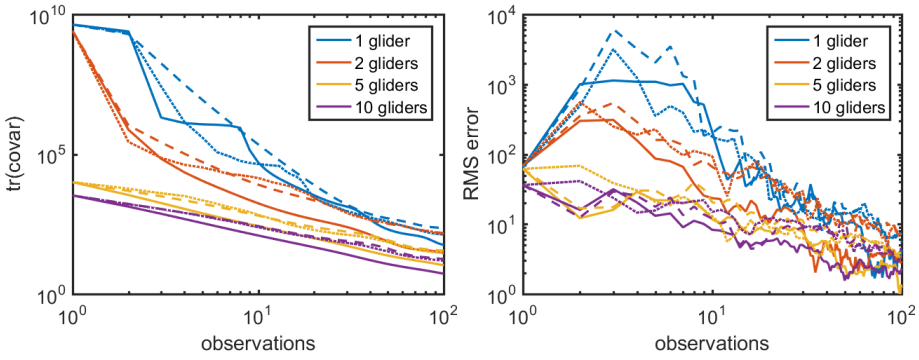

(a) Covariance trace, flow about center. (b) RMS error, flow about center.

Figure 2: (A) Decrease in covariance trace with number of observations for flow about center. For each glider cohort (blue, red, yellow, and purple curves respectively represent 1-, 2-, 5-, and 10-glider cohorts), the solid, dashed, and dotted curves respectively represent locally optimal control, no control, and random control. Optimal control performs systematically better, particularly as the number of observations increases. (B) Same as in (A), but with RMS error used as the performance metric. Initial growth in the RMS error for the 1- and 2-glider cohorts reflects the high impact of noisy observations on these underresolved cases.

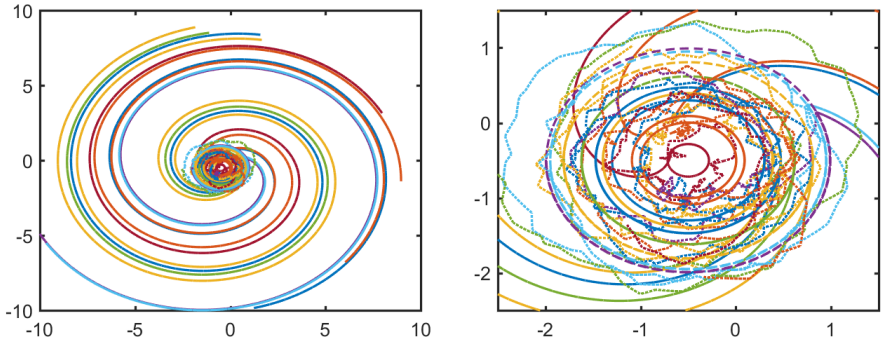

(a) Optimal paths, flow about center. (b) Blowup of (A).

Figure 3: (A) Optimally controlled glider velocities for flow about center. Solid lines are optimally controlled, dashed lines have no control, and dotted lines have random control. (B) Same as in (A) but with axes rescaled to give better picture of uncontrolled and randomly controlled paths.

there is little systematic improvement in the estimate quality for the first few iterations since the estimate is unable to accurately inform the optimization

computation. Once a sufficiently reliable velocity estimate is established, however, the improvement is substantial. This is particularly striking in the case of the single-glider cohort, where the RMS error drops precipitously between the 60th and 70th observation, breaking away from its uncontrolled counterpart. Also included as dotted lines in Fig. 2 are the assimilation results using randomly controlled trajectories, where the direction of control is drawn from a uniform distribution of angles on $[0, 2\pi)$ and held constant between observations. The randomly controlled trajectories are seen to offer some benefit in the intermediate observation times where the velocity estimate is poor, but they ultimately perform almost identically to the uncontrolled trajectories, and considerably worse than the optimally controlled trajectories.

To help understand the performance improvement offered by the optimally controlled gliders, their trajectories are plotted in Fig. 3 (A) and (B). As seen in Fig. (B), the gliders simply orbit around the center in the absence of control, with similar albeit erratic trajectories in the case of random control. The optimal control, however, pushes the gliders systematically toward larger distances from the origin. This is due to the fact that the information provided by observation operator 34 increases with larger values of x and y. Optimization of the expected posterior covariance trace naturally pushes the gliders towards these values that provide the highest information gain.

4.2. Flow about an unstable node

The simulations performed around an unstable node use a velocity field given by

$$v_0 = \begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, \qquad A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \tag{37}$$

such that the center is located at $(-1/2,\, 1/2)$. Figure 4 again shows the advantage of using optimal control, which in this case shows most clearly at intermediate stages of the assimilation process, after sufficient information has been acquired about the flow to produce effective control, and before sufficient observations have been taken to render all methods equally effective. The exception to this is the case of a single glider, where uncontrolled observations do not successfully learn the flow within the 100 observations taken. The paths taken by these gliders are shown in Fig. 5, where the most salient point is that the controlled gliders follow trajectories that asymptote toward radial lines emanating from the origin for the reason discussed above that this optimizes information gain from observation operator 34. The uncontrolled and randomly controlled paths asymptote toward radial lines emanating from the fixed point $(-1/2,\, 1/2)$, with the difference between these asymptotic behaviors most noticeable in Fig. 5b.

4.3. Flow about a saddle

The simulations performed around a saddle use a velocity field given by

$$v_0 = \begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, \qquad A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{38}$$

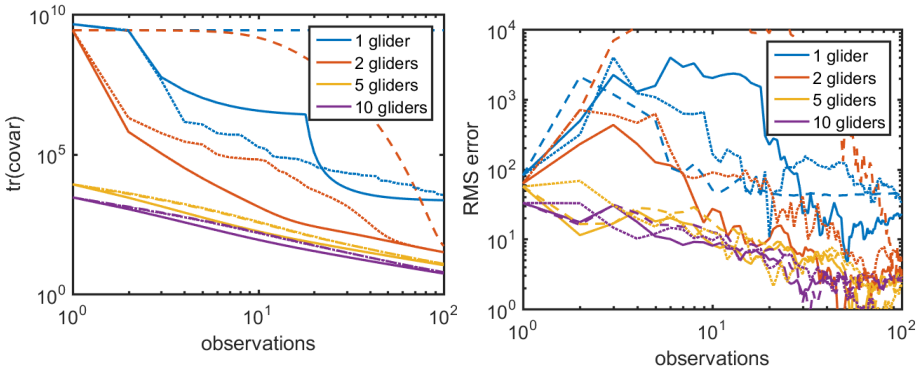

(a) Covariance trace, flow about unstable node.

(b) RMS error, flow about unstable node.

Figure 4: (A) Decrease in covariance trace with number of observations for flow about unstable node. For each glider cohort (blue, red, yellow, and purple curves respectively represent 1-, 2-, 5-, and 10-glider cohorts), the solid, dashed, and dotted curves respectively represent locally optimal control, no control, and random control. Optimal control performs systematically better for sufficiently large number of observations. (B) Same as in (A), but with RMS error used as the performance metric. Initial growth in the RMS error for the 1- and 2-glider cohorts reflects the high impact of noisy observations on these underresolved cases.

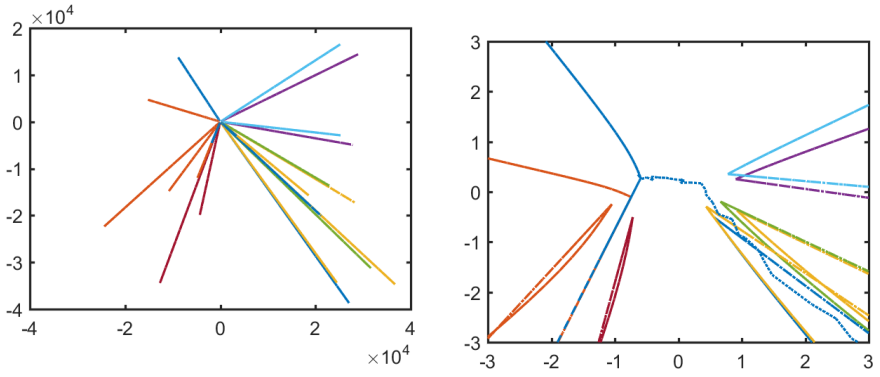

(a) Optimal paths, flow about unstable node.

(b) Blowup of (A).

Figure 5: (A) Optimally controlled glider velocities for flow about unstable node. Solid lines are optimally controlled, dashed lines have no control, and dotted lines have random control. (B) Same as in (A) but with axes rescaled to give better picture of uncontrolled and randomly controlled paths.

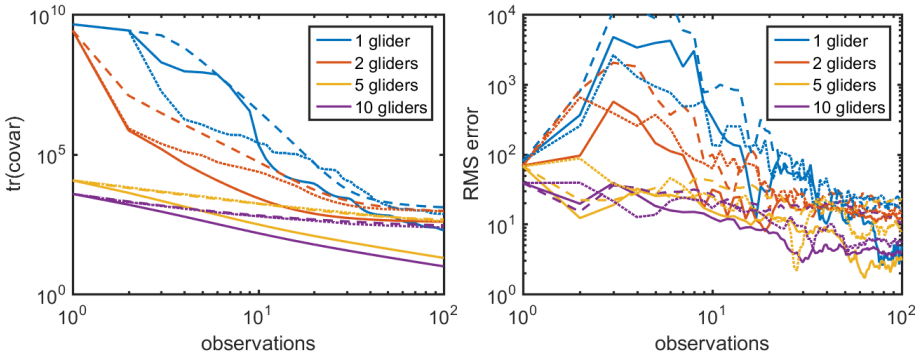

(a) Covariance trace, flow about saddle. (b) RMS error, flow about saddle.

Figure 6: (A) Decrease in covariance trace with number of observations for flow about saddle. For each glider cohort (blue, red, yellow, and purple curves respectively represent 1-, 2-, 5-, and 10-glider cohorts), the solid, dashed, and dotted curves respectively represent locally optimal control, no control, and random control. Optimal control performs systematically better. (B) Same as in (A), but with RMS error used as the performance metric. Initial growth in the RMS error for the 1- and 2-glider cohorts reflects the high impact of noisy observations on these underresolved cases.

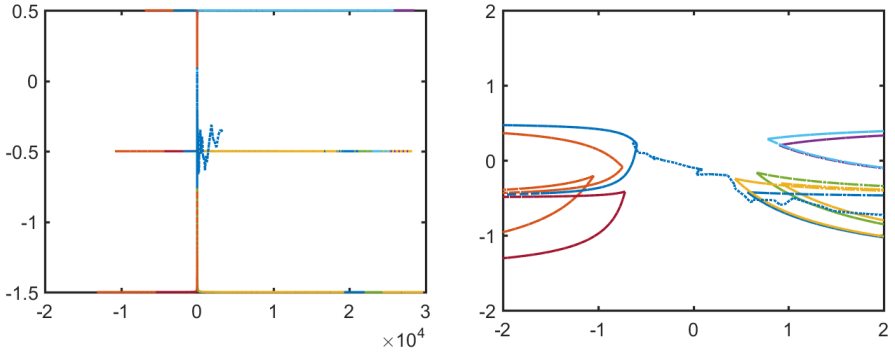

(a) Optimal paths, flow about saddle. (b) Blowup of (A).

Figure 7: (A) Optimally controlled glider velocities for flow about saddle. Solid lines are optimally controlled, dashed lines have no control, and dotted lines have random control. (B) Same as in (A) but with axes rescaled to give better picture of uncontrolled and randomly controlled paths.

such that the center is located at $(-1/2, -1/2)$. Figure 6 is similar to the cases before, showing systematic improvement in both covariance trace and RMS error when optimal control is used to guide the gliders. Their paths, shown

in Fig. 7, demonstrate that the optimal control orients itself transverse to the stable manifold in order to maximize the magnitude of the glider's y-value. The uncontrolled and randomly controlled gliders quickly collapse onto the unstable manifold $y = -1/2$.

4.4. Flow about a stable node

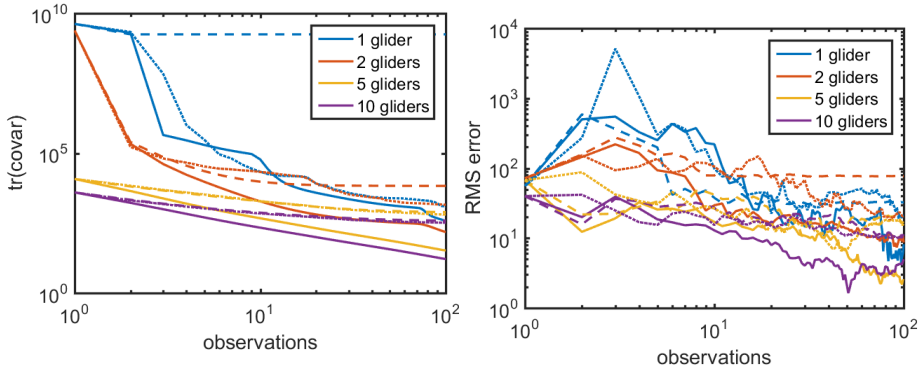

(a) Covariance trace, flow about stable node.

(b) RMS error, flow about stable node.

Figure 8: (A) Decrease in covariance trace with number of observations for flow about stable node. For each glider cohort (blue, red, yellow, and purple curves respectively represent 1-, 2-, 5-, and 10-glider cohorts), the solid, dashed, and dotted curves respectively represent locally optimal control, no control, and random control. Optimal control performs systematically better. (B) Same as in (A), but with RMS error used as the performance metric. Initial growth in the RMS error for the 1- and 2-glider cohorts reflects the high impact of noisy observations on these underresolved cases.

The simulations performed around a stable node use a velocity field given by

$$v_0 = \begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, \qquad A = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{39}$$

such that the center is located at $(1/2, -1/2)$. In this case, Figure 8 demonstrates how, in the absence of control, observations from the gliders are eventually rendered uninformative as the gliders collapse onto the fixed point. Random control offers some improvement, but optimal control is systematically better, particularly once sufficient information about the flow has been acquired. The paths shown in Fig. 10 demonstrate that the optimal control strives to keep the gliders from collapsing into the attractor.

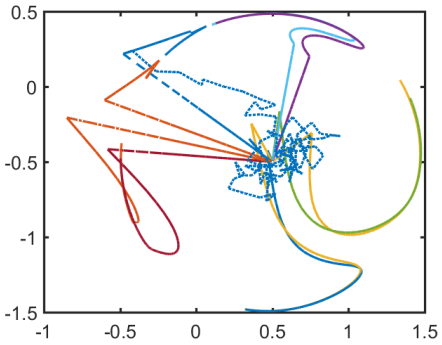

Figure 9: Optimal paths, flow about stable node.

Figure 10: (A) Optimally controlled glider velocities for flow about saddle. Solid lines are optimally controlled, dashed lines have no control, and dotted lines have random control. (B) Same as in (A) but with axes rescaled to give better picture of uncontrolled and randomly controlled paths.

5. Discussion

The methodology presented above provides a systematic way to iteratively improve the estimate of a velocity field obtained through the use of observations made by controlled autonomous vehicles. We have shown through the use of simple examples consisting of stationary linear flows near hyperbolic fixed points with different stability properties that this approach is effective, providing significant improvement over uncontrolled AVs and over AVs that are driven at maximal speed in random directions between each assimilation stage. We obtain the locally optimal paths through solution of the Euler-Lagrange equations resulting from an objective function provided by the expected variance of the posterior distribution which, in the case of direct velocity observations, depends on the locations of the AVs at each assimilation time. As is often the case in optimal control problems, these equations are not always solvable, and we provide a regularization method that provably works in the case of trivial flows and has resolved all cases of nonexistence of solutions in the linear flows considered here.

The example flows considered here are simple linear flows for the purpose of clear illustration of the optimization results. We now consider the natural extensions to this method towards application to more realistic flows, all of which are the subject of ongoing work.

5.1. Nonlinear, random, and time-dependent flows

A primary motivation for the work described in [3] was the observation that Lagrangian *drifters* (i.e., with no capacity for self-propulsion) often encounter

flow barriers that prevent their effectiveness in learning regional flows. The kinematic, incompressible flow [50] used as a test case in that work took the form of a stream function given by

$$\psi = -\pi y + \sin(2\pi x)\sin(2\pi y), \quad (x, y) \in [0, 1] \times [0, 0.5], \tag{40}$$

with associated (stationary) velocity field

$$v_1 = -\frac{\partial \psi}{\partial y} = \pi - 2\pi \sin(2\pi x)\cos(2\pi y) \tag{41}$$

$$v_2 = \frac{\partial \psi}{\partial x} = 2\pi \cos(2\pi x)\sin(2\pi y). \tag{42}$$

If the mean flow $\langle v \rangle = (\pi, 0)^T$ is assumed known with the rest assumed to satisfy periodic boundary conditions in x and Dirichlet boundary conditions in y, an estimate $\tilde{\psi}$ of the stream function can be represented through a finite Fourier basis given by

$$\tilde{\psi} = \pi y + \sum_{n=1}^{N} a_{on} \sin(2n\pi y) + \sum_{m=1}^{M}\sum_{n=1}^{N} [a_{mn}\cos(2m\pi x) + b_{mn}\sin(2m\pi x)]\sin(2n\pi y). \tag{43}$$

Expressing the state vector $\hat{\psi}$ as a column of these coefficients with dimension $N(1 + 2M)$, the observation operator mapping from state space to observation space then takes the form

$$H = \begin{pmatrix} k_1\cos(k_1 y_1) & \dots & k_1\cos(\mu_1 x_1)\cos(k_1 y_1) & k_1\sin(\mu_1 x_1)\cos(k_1 y_1) & \dots \\ 0 & \dots & \mu_1\sin(\mu_1 x_1)\sin(k_1 y_1) & -\mu_1\cos(\mu_1 x_1)\sin(k_1 y_1) & \dots \\ k_1\cos(k_1 y_2) & \dots & & & \\ \vdots & & & & \end{pmatrix}, \tag{44}$$

from which one can see the dependence on glider positions $z_k = (x_k, y_k)^T$. More complex flows imply more significant difficulties with nonexistence of solutions to the Euler-Lagrange equations, but preliminary results are nevertheless promising. Depicted in Fig. 11 are comparisons of controlled and uncontrolled paths, with their related reduction in the spatial variance of the stream function estimate. Whereas the controlled path is trapped within the gyre, producing observations that do little to reduce uncertainty outside of the gyre, the optimal control algorithm produces a path that crosses the separatrix into the meandering flow, producing a corresponding reduction in uncertainty there. This work is ongoing, as is the incorporation of weakly nonlinear time evolution models for the flow that fit naturally into the framework of the extended Kalman filter.

5.2. Observations of position

As described in Sec. 2, direct observations of the velocity local to a glider produce an observation operator of the form given in Eqn. 4, where it is clear how the glider positions at assimilation time influence the Kalman gain and

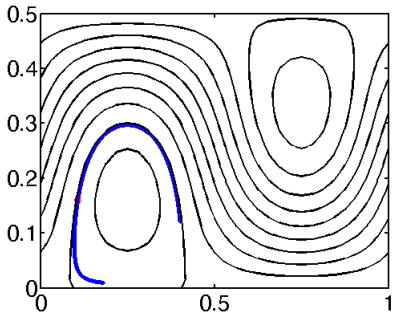

(a) Uncontrolled path, double gyre.

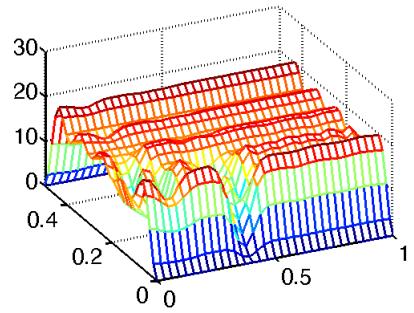

(b) Variance, uncontrolled path in double gyre.

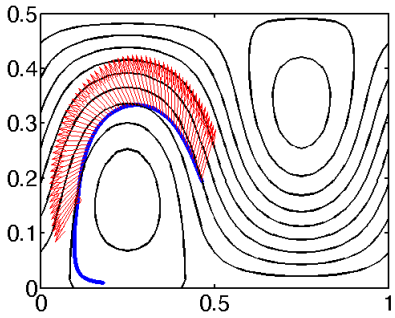

(c) Controlled path, double gyre.

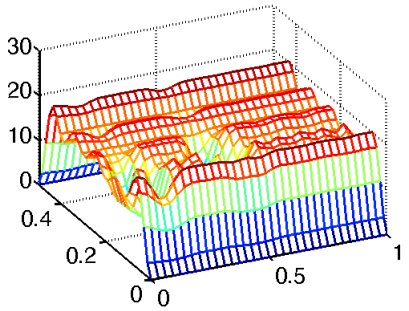

(d) Variance, controlled path in double gyre.

Figure 11: (A) Typical uncontrolled path (blue) trapped in gyre over two evolution periods with an observation (red circle) between them. Black curves are streamlines of true flow. (B) Variance in estimate $\tilde{\psi}(x, y)$ after two observations along uncontrolled path from (A). (C) Typical controlled path (blue), with control vector depicted in red. (D) Variance in $\tilde{\psi}(x, y)$ after two observations along controlled path from (C).

therefore the expected posterior covariance. When it is the gliders' positions that are observed, the methodology proposed in Ref. ([43]) poses as the state the augmented vector consisting of a parametric representation of the velocity field concatenated with the glider positions themselves. This produces the observation operator given in Eqn. 5, where it is no longer immediately clear how to formulate the dependence of expected posterior variance on glider path in such a way as to compute an optimal control. In this case, since the estimated state includes the glider positions themselves, the covariance of the state must be evolved between assimilation times, and this time evolution of the covariance is where the control plays a role. The extension of our control algorithm to this case will be presented elsewhere.

[1] A. R. Robinson, D. B. Haidvogel, Dynamical Forceast Experiments with a Baroptropic Open Ocean Model, Journal of Physical Oceanography 10 (1981) 1928.

[2] N. Chen, A. J. Majda, X. T. Tong, Information barriers for noisy Lagrangian tracers in filtering random incompressible flows, Nonlinearity 27 (9) (2014) 2133–2163. doi:10.1088/0951-7715/27/9/2133.

[3] D. McDougall, C. K. R. T. Jones, Decreasing flow uncertainty in bayesian inverse problems through lagrangian drifter control.

[4] E. Kalnay, Atmospheric modeling, data assimilation and predictability, Cambridge University Press, 2002.

[5] R. E. Kalman, A New Approach to Linear Filtering and Prediction Problems, Journal of Basic Engineering 82 (Series D) (1960) 35–45.

[6] R. E. Kalman, R. S. Bucy, New results in linear filtering and prediction theory, Journal of Basic Engineering 83 (1961) 95–107.

[7] H. W. Sorenson, Kalman filtering: theory and application, IEEE, 1960.

[8] G. Evensen, Data Assimilation: The Ensemble Kalman Filter, Springer, 2006.

[9] P. L. Houtekamer, H. L. Mitchell, Data Assimilation Using an Ensemble Kalman Filter Technique, Monthly Weather Review 126 (1998) 796–811.

[10] J. L. Anderson, A Local Least Squares Framework for Ensemble Filtering, Monthly Weather Review 131 (4) (2003) 634–642.

[11] A. C. Lorenc, S. P. Ballard, R. S. Bell, N. B. Ingleby, P. L. F. Andrews, D. M. Barker, J. R. Bray, A. M. Clayton, T. Dalby, D. Li, T. J. Payne, F. W. Saunders, The Met. Office global three-dimensional variational data assimilation scheme, Quarterly Journal of the Royal Meteorological Society 126 (570) (2000) 2991–3012. doi:10.1002/qj.49712657002.

[12] L. Bengtsson, 4-dimensional assimilation of meteorological observations, World Meteorological Organization.

[13] J. M. Lewis, J. C. Derber, The use of adjoint equations to solve a variational adjustment problem with advective constraints, Tellus A (37A) (1985) 309–322.

[14] A. C. Lorenc, Analysis methods for numerical weather prediction, Quarterly Journal of the Royal Meteorological Society 112 (1986) 1177–1194.

[15] F.-X. Le Dimet, O. Talagrand, Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects, Tellus A (38A) (1986) 97–110.

[16] O. Talagrand, P. Courtier, Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory, Quarterly Journal of the Royal Meteorological Society 113 (1987) 1311–1328.

[17] P. Courtier, J.-N. Thépaut, A. Hollingsworth, A strategy for operational implementation of 4D-Var, using an incremental approach, Quarterly Journal of the Royal Meteorological Society 120 (519) (1994) 1367–1387.

[18] A. S. Lawless, S. Gratton, N. K. Nichols, An investigation of incremental 4D-Var using non-tangent linear models, Tech. rep. (2005).

[19] A. S. Lawless, S. Gratton, N. K. Nichols, Approximate iterative methods for variational data assimilation, Tech. Rep. April (2005).

[20] A. Doucet, N. de Freitas, N. Gordon, Sequential Monte Carlo Methods In Practice, 2001.

[21] P. J. V. Leeuwen, Nonlinear data assimilation in geosciences: an extremely efficient particle filter, Quarterly Journal of the Royal Meteorological Society 136 (653) (2010) 1991–1999. doi:10.1002/qj.699.

[22] P. Bickel, B. Li, T. Bengtsson, Sharp failure rates for the bootstrap particle filter in high dimensions, in: Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh, Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2008, pp. 318–329.

[23] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, Equation of State Calculations by Fast Computing Machines, The Journal of Chemical Physics 21 (6) (1953) 1087. doi:10.1063/1.1699114.

[24] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, Biometrika 57 (1) (1970) 97–109. doi:10.1093/biomet/57.1.97.

[25] S. L. Cotter, G. O. Roberts, A. Stuart, D. White, et al., MCMC methods for functions: modifying old algorithms to make them faster, Statistical Science 28 (3) (2013) 424–446.

[26] S. L. Cotter, M. Dashti, J. C. Robinson, A. M. Stuart, Bayesian inverse problems for functions and applications to fluid mechanics, Inverse Problems 25 (11) (2009) 115008. doi:10.1088/0266-5611/25/11/115008.

[27] S. L. Cotter, M. Dashti, A. M. Stuart, Approximation of Bayesian inverse problems for PDEs, SIAM Journal of Numerical Analysis 48 (1) (2010) 322–345.

[28] W. Lee, D. McDougall, A. M. Stuart, Kalman filtering and smoothing for linear wave equations with model error, Inverse Problems 27 (9) (2011) 095008. doi:10.1088/0266-5611/27/9/095008.

[29] A. Apte, C. K. R. T. Jones, A. M. Stuart, J. Voss, Data assimilation: Mathematical and statistical perspectives, International Journal for Numerical Methods in Fluids 56 (2008) 1033–1046. doi:10.1002/fld.

[30] A. Apte, M. Hairer, A. M. Stuart, J. Voss, Sampling the posterior: An approach to non-Gaussian data assimilation, Physica D: Nonlinear Phenomena 230 (1-2) (2007) 50–64. doi:10.1016/j.physd.2006.06.009.

[31] A. Apte, C. K. R. T. Jones, A. M. Stuart, A Bayesian approach to Lagrangian data assimilation, Tellus A 60 (2) (2008) 336–347. doi:10.1111/j.1600-0870.2007.00295.x.

[32] R. Herbei, I. McKeague, Hybrid Samplers for Ill-Posed Inverse Problems, Scandinavian Journal of Statistics 36 (4) (2009) 839—-853.

[33] J. P. Kaipio, V. Kolehmainen, E. Somersalo, M. Vauhkonen, Statistical inversion and Monte Carlo sampling methods in electrical impedance tomography, Inverse problems 16 (5) (2000) 1487.

[34] K. Mosegaard, A. Tarantola, Monte Carlo sampling of solutions to inverse problems, Journal of Geophysical Research 100 (B7) (1995) 12431—-12447. doi:10.1029/94JB03097.

[35] G. O. Roberts, Weak convergence and optimal scaling of random walk Metropolis Algorithms, Annals of Applied Probability 7 (1) (1997) 110–120.

[36] G. O. Roberts, J. S. Rosenthal, Optimal scaling of discrete approximations to Langevin diffusions, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 60 (1) (1998) 255–268. doi:10.1111/1467-9868.00123.

[37] G. O. Roberts, J. S. Rosenthal, Optimal scaling for various Metropolis-Hastings algorithms, Statistical Science 16 (4) (2001) 351–367. doi:10.1214/ss/1015346320.

[38] A. Beskos, G. O. Roberts, A. M. Stuart, Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions, The Annals of Applied Probability 19 (3) (2009) 863–898. doi:10.1214/08-AAP563.

[39] Y. F. Atchadé, J. S. Rosenthal, On adaptive Markov chain Monte Carlo algorithms, Bernoulli 11 (5) (2005) 815–828. doi:10.3150/bj/1130077595.

[40] Y. F. Atchadé, An Adaptive Version for the Metropolis Adjusted Langevin Algorithm with a Truncated Drift, Methodology and Computing in Applied Probability 8 (2) (2006) 235–254. doi:10.1007/s11009-006-8550-0.

[41] A. J. Mariano, A. Griffa, T. M. zgkmen, E. Zambianchi, Lagrangian analysis and predictability of coastal and ocean dynamics 2000, Journal of Atmospheric and Oceanic Technology 19 (7) (2002) 1114–1126.

[42] N. E. Leonard, D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, F. Zhang, Coordinated Control of an Underwater Glider Fleet in an Adaptive Ocean Sampling Field Experiment in Monterey Bay, Journal of Field Robotics 27 (6) (2010) 718–740. doi:10.1002/rob.

[43] L. Kuznetsov, K. Ide, C. K. R. T. Jones, A Method for Assimilation of Lagrangian Data, Monthly Weather Review 131 (2003) 2247–2260.

[44] R. S. Liptser, A. N. Shiryaev, Statistics of Random Processes I. General Theory, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.

[45] L. I. Piterbarg, Optimal estimation of Eulerian velocity field given Lagrangian observations, Applied Mathematical Modelling 32 (10) (2008) 2133–2148. doi:10.1016/j.apm.2007.07.011.

[46] A. M. Stuart, Inverse problems: a bayesian perspective, Acta Numerica 19 (2010) 451–559.

[47] K. Ide, P. Courtier, M. Ghil, A. C. Lorenc, Uni ed notation for data assimilation: operational, sequential and variational, J. Meteorol. Soc. Japan (1997) 181189.

[48] A. E. Bryson Jr., Y.-C. Ho, Applied Optimal Control: Optimization, estimation, and control, Hemisphere Publishing Corporation, 1975.

[49] H. B. Keller, Numerical solution of two point boundary value problems, 4th Edition, no. 24 in CBMS-NSF regional conference series in applied mathematics, Society for Industrial and Applied Mathematics, Philadelphia, Pa, 1990.

[50] R. M. Samelson, S. Wiggins, Lagrangian transport in geophysical jets and waves the dynamical systems approach, Springer, New York, N.Y., 2006.